EXHIBIT L

Exhibit 4 – SYNAPSE-1

'156 Patent

Claim Limitation (Claim 7)	Exemplary Disclosure
[**156a**] A device comprising:	SYNAPSE-1 discloses a device. Specifically, the SYNAPSE -1 is a neurocomputer. *See, e.g.*: "A general purpose neurocomputer, SYNAPSE-1, is presented which exhibits a multi processor and memory architecture. It offers wide flexibility with respect to neural algorithms and a speed-up factor of several orders of magnitude—including learning. The computational power is provided by a 2-dimensional systolic array of neural signal processors. Since the weights are stored outside these NSPs memory size and processing power can be adapted individually to the applicational needs. A neural Algorithms Programming Language, embedded in C++, has been defined for the user to cope with the neurocomputer. In a benchmark test the prototype of SYNAPSE-I was 8000 times as fast as a standard workstation." U. Ramacher, et al., *Multiprocessor and Memory Architecture of the Neurocomputer SYNAPSE-1*, INT. J. NEURAL SYST., 1993, at 1034, https://pubmed.ncbi.nlm.nih.gov/8049796/ "SYNAPSE–1 is a modular system whose building blocks are arranged in a two–dimensional structure [3, 4]. The building blocks are (Fig. 1) a two-dimensional array of neuro signal processors MA16, weight memories, data units, and a control unit. The central part of the neurocomputer is the matrix of processing elements. The processing elements receive data from data units at the left edge of the processing array. The synaptic weights that are required for processing are input from the weight memories at the top edge. Partial results are computed and pipelined along the rows of the matrix to the right. After the matrix-vector products have been computed the results are routed back to the Data Units. Here all computational steps are executed that cannot be done in the processing array itself. Processing is done under central control by the Control Unit." U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	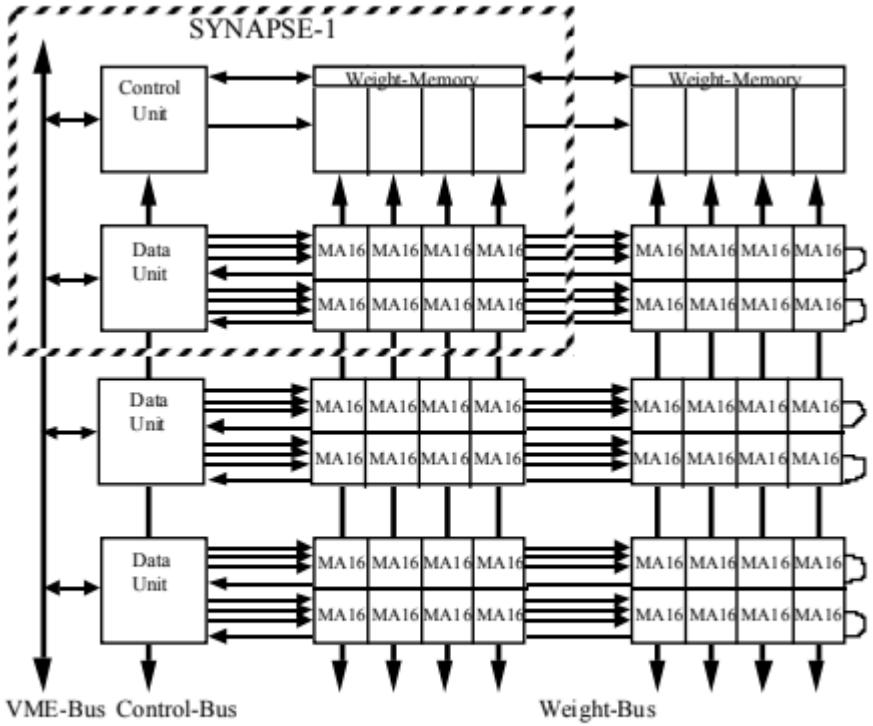 Figure 1: System architecture. U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.
[**156b**] at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SYNAPSE-1 discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See, e.g.*: "The compute-intensive operations (e.g. matrix multiplication) common to all neural

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	algorithms including learning paradigms are executed on an array of neural signal processors MA16, each providing a peak performance of $640 \cdot 10^6$ connections (16x16 bit) per second at 40 MHz [2]. The non-compute-intensive operations are performed on the 'Data Unit', while the 'Control Unit' cares for the concatenation of compute-intensive and non-compute-intensive operations as well as the communication with the host." U. Ramacher, et al., *Multiprocessor and Memory Architecture of the Neurocomputer SYNAPSE-1*, INT. J. NEURAL SYST., 1993, at 1035, https://pubmed.ncbi.nlm.nih.gov/8049796/ "The processor board hosts a 2-dimensional systolic array of neural signal processors, arranged in two rows by four columns. Each MA16 is accompanied by a local memory for intermediate data (Z-Memory). Both processor rows are connected to the same weights bus, thus executing the same neural net operations for different input patterns. The MA16s of a row form a linear systolic array where input data as well as (intermediate) results propagate from the Data Unit through the MA16s back to the Data Unit. Each MA16 itself contains a linear systolic array of 4 processing modules PM (fig. 2). The data paths and operations of a PM are set up such that both neural algorithms and classical signal processing algorithms (e.g. DCT, correlation, etc.) can be performed." U. Ramacher, et al., *Multiprocessor and Memory Architecture of the Neurocomputer SYNAPSE-1*, INT. J. NEURAL SYST., 1993, at 1036, https://pubmed.ncbi.nlm.nih.gov/8049796/ "SYNAPSE–1 is a modular system whose building blocks are arranged in a two–dimensional structure [3, 4]. The building blocks are (Fig. 1) a two-dimensional array of neuro signal processors MA16, weight memories, data units, and a control unit. The central part of the neurocomputer is the matrix of processing elements. The processing elements receive data from data units at the left edge of the processing array. The synaptic weights that are required for processing are input from the weight memories at the top edge. Partial results are computed and pipelined along the rows of the matrix to the right. After the matrix-vector products have been computed the results are routed back to the Data Units. Here all computational steps are executed that cannot be done in the processing array itself. Processing is done under central control by the Control Unit." U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*,

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.

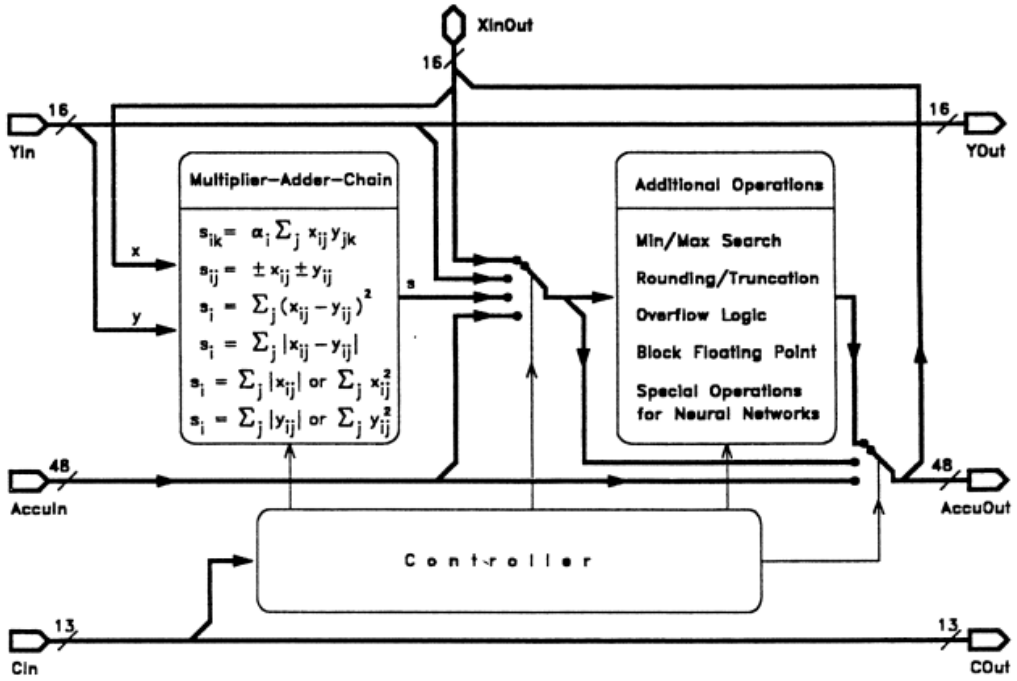

Figure 2 Block diagram of a processor module PM

U. Ramacher, et al., *Multiprocessor and Memory Architecture of the Neurocomputer SYNAPSE-1*, INT. J. NEURAL SYST., 1993, at 1036, https://pubmed.ncbi.nlm.nih.gov/8049796/

"The Neuroprocessor Unit executes the most time critical computations required for any simulation of neural networks including learning, e.g., the computation of matrix–vector products. It consists of a two–dimensional systolic array of neural signal processors MA16 [2], arranged in two rows by four columns (Fig. 2). For large matrix–vector products the MA16 processors calculate the results in several passes. In addition to the data and weight

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	ports, each chip there- fore provides input/output ports for partially accumulated results. These ports are horizontally connected between adjacent MA16 processors. In this way a row of MA16s form a linear systolic array where input data as well as partial results propagate from the Data Unit through the MA16s back to the Data Unit." U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.
[**156c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and	As explained below and in the Responsive Contentions Regarding Non-Infringement and Invalidity ("Responsive Contentions"), it would have been obvious to one of skill in the art based on the disclosures in SYNAPSE-1 (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32, or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. Reduced precision computations were commonly used in connection with neurocomputers like the SYNAPSE-1. *See, e.g.*: "Other systems try to take advantage of other peculiarities of ANN algorithms, such as a reduced precision required in the computations. This makes it possible to develop ad-hoc processing elements which are characterized by a small size and cost. In turn, this often enables one to design systems with a very high processing element count and therefore to increase the degree of parallelism." Paolo Ienne, *Digital Systems for Neural Networks*, PROC. SPIE 10279, Digital Signal Processing Technology: A Critical Review, April 25, 1995, at 11.

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	"BP implementations typically use 32-bit floating point math. This largely eliminates scaling, precision and dynamic range issues. Efficient hardware implementation dictates integer arithmetic units with precision no greater than required. Baker [Bak90] has shown 16-bit integer weights are sufficient for BP training and much lower values adequate for use after training." Hal McCartor, *Back Propagation Implementation on the Adaptive Solutions CNAPS Neurocomputer Chip*, ADVANCES IN NEURAL INFORMATION PROCESSING SYSTEMS 3, 1990, at 1029. SYNAPSE-1 discloses the use of arithmetic units designed to perform calculations using 16-bit fixed-point math for calculations that typically use 32-bit floating point math. *See, e.g.*: "Each MA16 contains four systolic chains of four processing elements. A single MA16 processor, operating at 40 MHz, represents a computing power of 640 · 106 connections/s. Products are computed as 16 × 16 bits fixed point integers whereas the results are summed up in 48 bits. With a 2×4 array of neuro signal processors MA16, a single Neuroprocessor Unit can compute up to ≈ 5 · 109 connections per second." Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995. For the reasons explained in the Responsive Contentions, it would have been obvious to one of skill in the art to have substituted the fixed-point number format used in SYNAPSE-1 for a floating-point format that met the claimed minimum range and precision requirements, and to have used the reduced-precision floating-point number formats disclosed in Tong, Dockser, Belanovic / Belanovic and Leeser, Shirazi, Sudha, Aty, and TMS 320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger, either alone or in combination. *See also* Appendix to Responsive Contentions (detailing error rates associated with different mantissa sizes).
[**156d**] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit;	SYNAPSE-1 discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. *See, e.g.*:

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	"The Control Unit is divided into 2 main groups: A MC68040 CPU takes over the communication with the host and the Data Unit CPU. It also controls the other parts of the board, thereby synthesizing the complete algorithm to be performed (SYNAPSE = synthesis of neural algorithms on a parallel systolic engine). A microprogram sequencer controls the MA16 array along with all components which must operate synchronously to the array. These are the W-Memory, the Z- Memory, FIFOs, and the Backplane Transceivers. Therefore a microprogram generally corresponds to the compute-intensive part of an elementary operation of SYNAPSE-1." U. Ramacher, et al., *Multiprocessor and Memory Architecture of the Neurocomputer SYNAPSE-1*, INT. J. NEURAL SYST., 1993, at 1037, https://pubmed.ncbi.nlm.nih.gov/8049796/

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	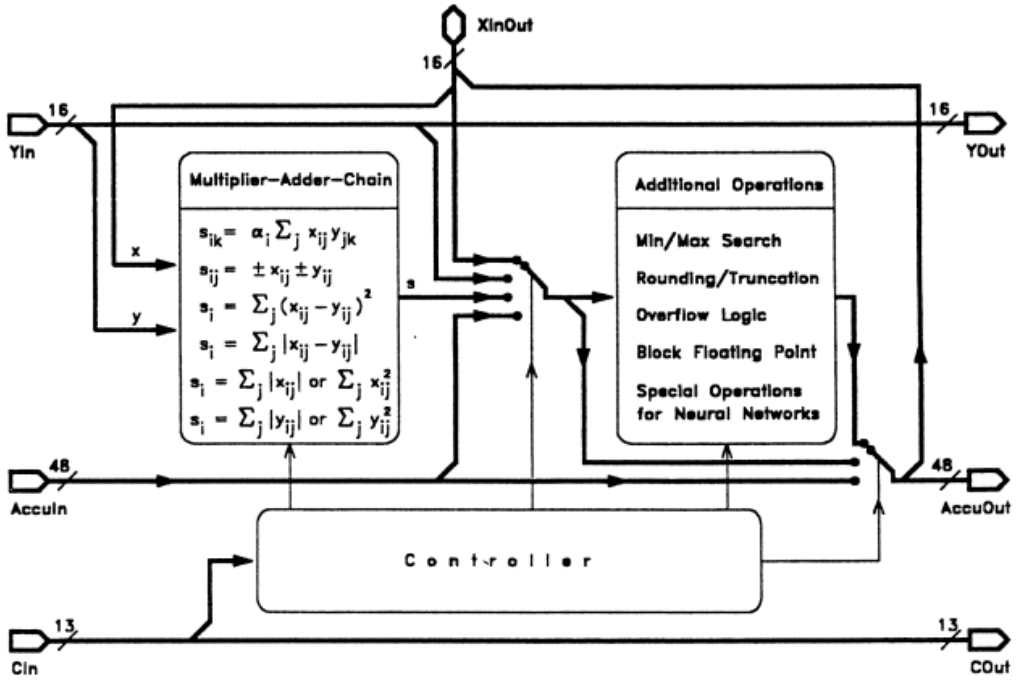 Figure 2 Block diagram of a processor module PM U. Ramacher, et al., *Multiprocessor and Memory Architecture of the Neurocomputer SYNAPSE-1*, INT. J. NEURAL SYST., 1993, at 1036, https://pubmed.ncbi.nlm.nih.gov/8049796/ "The Control Unit has two main responsibilities. First it coordinates all activities of SYNAPSE–1. This is done by a programmable sequencer that generates appropriate control signals and addresses for the other units. Second it provides the communication of the neurocomputer with a host system, usually a workstation used for programming and as the user interface." U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	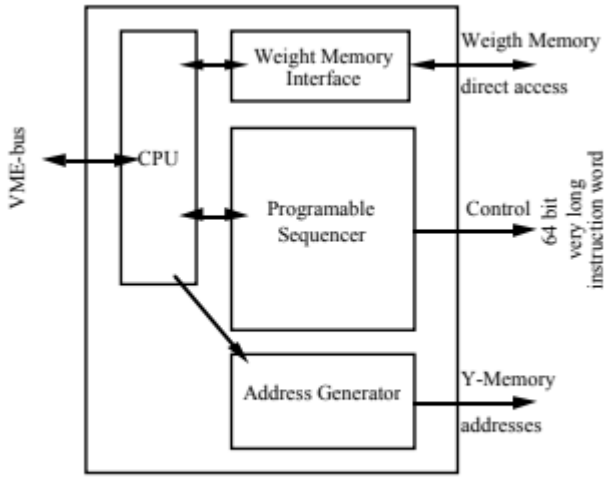 Figure 5: Control Unit. U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.
[**156e**] wherein the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine;	SYNAPSE-1 discloses at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine. *See, e.g.*: "The Control Unit is divided into 2 main groups: A MC68040 CPU takes over the communication with the host and the Data Unit CPU. It also controls the other parts of the board, thereby synthesizing the complete algorithm to be performed (SYNAPSE = synthesis of neural algorithms on a parallel systolic engine).

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	A microprogram sequencer controls the MA16 array along with all components which must operate synchronously to the array. These are the W-Memory, the Z- Memory, FIFOs, and the Backplane Transceivers. Therefore a microprogram generally corresponds to the compute-intensive part of an elementary operation of SYNAPSE-1." U. Ramacher, et al., *Multiprocessor and Memory Architecture of the Neurocomputer SYNAPSE-1*, INT. J. NEURAL SYST., 1993, at 1037, https://pubmed.ncbi.nlm.nih.gov/8049796/ "The Control Unit has two main responsibilities. First it coordinates all activities of SYNAPSE–1. This is done by a programmable sequencer that generates appropriate control signals and addresses for the other units. Second it provides the communication of the neurocomputer with a host system, usually a workstation used for programming and as the user interface." U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995. 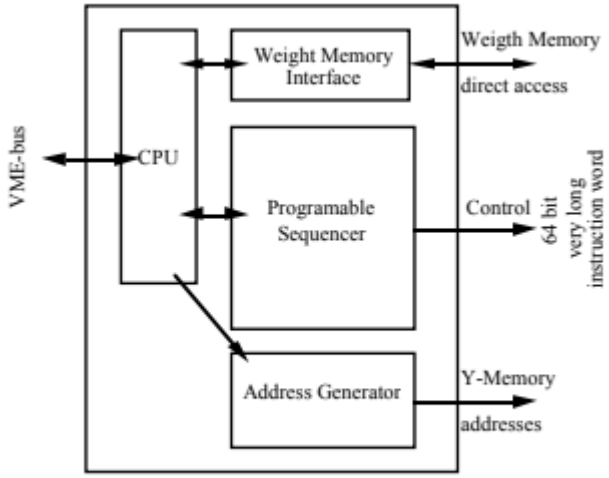 Figure 5: Control Unit. U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel*

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	Neurocomputer System, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.
[**156f**] and, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	SYNAPSE-1 discloses the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. The SYNAPSE-1 includes Neuroprocessor Units containing eight MA16 chips, each of which contains four processing elements. "The Neuroprocessor Unit executes the most time critical computations required for any simulation of neural networks including learning, e.g., the computation of matrix–vector products. It consists of a two–dimensional systolic array of neural signal processors MA16 [2], arranged in two rows by four columns (Fig. 2). For large matrix–vector products the MA16 processors calculate the results in several passes. In addition to the data and weight ports, each chip therefore provides input/output ports for partially accumulated results. These ports are horizontally connected between adjacent MA16 processors. In this way a row of MA16s form a linear systolic array where input data as well as partial results propagate from the Data Unit through the MA16s back to the Data Unit. Each MA16 contains four systolic chains of four processing elements." U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995. SYNAPSE-1 is modular, so additional Neuroprocessor Units can be added to the SYNAPSE-1 to achieve higher data parallelism. "The processing power and the storage capacity of the memories can be adapted to the application needs. The minimal system configuration, which is operating since autumn 1992, consists of one Neuroprocessor Unit, one Weight Memory Unit, one Data Unit, and one Control Unit (see marked box in Fig. 1). For applications needing a higher input data rate the systolic array can be extended using more rows of Data Units and Neuroprocessor

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 7)	Exemplary Disclosure
	Units. In this way higher data parallelism is obtained by redistributing the same data over more processors operating with the same weights. For applications that require larger neural networks, i.e., more processing power as well as more weight memory space, the systolic array can be extended using more columns of Weight Memory Units and Neuroprocessor Units. System extension is easily achieved by using the appropriate number of boards in a system crate. The interconnection is provided by two special–purpose buses for weights and control signals, and a system bus. The VMEbus has been chosen as the system bus for simple interfacing to off–the– shelf hardware like the host computer or I/O devices." U. Ramacher, et al., *SYNAPSE–1: A High–Speed General Purpose Parallel Neurocomputer System*, Proceedings of 9th International Parallel Processing Symposium, Santa Barbara, CA, 1995.

Exhibit 4 – SYNAPSE-1

'273 Patent

Claim Limitation (Claim 53)	Exemplary Disclosure
[**273a**] A device:	SYNAPSE-1 discloses a device. Specifically, the SYNAPSE -1 is a neurocomputer. *See* [**156a**].
[**273b**] comprising at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SYNAPSE-1 discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**273c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input;	For the reasons explained above and in the Responsive Contentions, it would have been obvious to one of skill in the art based on the disclosures in CNAPS (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**].
[**273d**] wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	SYNAPSE-1 discloses the number of LPHDR execution units in the device exceeds by at least one hundred SYNAPSE-1 the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. *See* [**156f**].

Exhibit 4 – SYNAPSE-1

'961 Patent

Claim Limitation (Claim 4)	Exemplary Disclosure
[**961a**] A device comprising:	SYNAPSE-1 discloses a device. Specifically, the SYNAPSE -1 is a neurocomputer. *See* [**156a**].
[**961b**] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SYNAPSE-1 discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**961c**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and	For the reasons explained above and in the Responsive Contentions, it would have been obvious to one of skill in the art based on the disclosures in CNAPS (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**].
[**961d**] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.	SYNAPSE-1 discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. *See* [**156d**].

Exhibit 4 – SYNAPSE-1

Claim Limitation (Claim 13)	Exemplary Disclosure
[**961e**] A device comprising:	SYNAPSE-1 discloses a device. Specifically, the SYNAPSE -1 is a neurocomputer. *See* [**156a**].
[**961f**] a plurality of components comprising:	SYNAPSE-1 discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* **[156b].** SYNAPSE-1 discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. *See above* [**156d**].
[**961g**] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	SYNAPSE-1 discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. *See* [**156b**].
[**961h**] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.	For the reasons explained above and in the Responsive Contentions, it would have been obvious to one of skill in the art based on the disclosures in CNAPS (alone or in combination with the reduced precision floating point teachings of Dockser, Tong, Belanovic / Belanovic and Leeser, Lee, Shirazi, Aty, Sudha, and TMS320C32 or the logarithmic format disclosed in GRAPE-3 and Hoefflinger) that the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. *See* [**156c**].